

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: MULTIPLE STORAGE ARRAY CONTROL
APPLICANTS: JOSEPH G. MURPHY AND ILYA LUBOVICH

CERTIFICATE OF MAILING BY EXPRESS MAIL

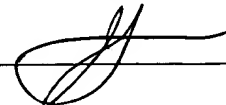
Express Mail Label No. EL228027128US

I hereby certify under 37 CFR §1.10 that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

November 1, 1999

Date of Deposit

Signature



Jason DeThomas

Typed or Printed Name of Person Signing Certificate

MULTIPLE STORAGE ARRAY CONTROL

Background of the Invention

5 The invention relates generally to the control of multiple storage devices, such as disk arrays.

Storage systems have grown enormously in both size and sophistication in recent years. These systems can typically include many large disk drive units controlled by
10 a complex, multi-tasking, disk drive controller such as the EMC Symmetrix disk drive controller. A large scale disk drive system can typically receive commands, such as I/O requests, from a number of host computers and can control a number of disk drive mass storage devices, each mass storage
15 device capable of storing in excess of several gigabits of data.

There is every reason to expect that both the sophistication and the size of the disk drive systems will continue to increase. As the systems increase in
20 complexity, so does a user's reliance upon the system for fast and reliable recovery and storage of data. Accordingly, the user typically uses data throughput and speed of response as a primary criteria for evaluating performance of the disk drive systems. As a result, mass
25 storage devices and the controllers which drive them have become quite sophisticated in trying to improve command response time. Systems such as the EMC Symmetrix disk drive controller system thus incorporate a large cache memory, and other techniques to improve the system throughput.

30 Of course, with the increase in size and complexity of such disk drive systems, comes a tremendous increase in the number of command requests making monitoring of the performance of the systems more difficult.

Summary of the Invention

The invention relates to a network architecture which facilitates the monitoring and control of a number of storage systems (e.g., arrays) by one or more clients.

5 In a general aspect of the invention, the network architecture includes at least one storage management server, connected between the storage systems and a number of clients, for providing operation status information associated with the storage systems to at least one of the
10 clients.

Among other advantages of the storage management server, is the ability for a client to access (e.g., via an Intranet) the storage network and check the status of the storage systems for any problems which might affect
15 throughput of data requested from a particular storage system. As the size of storage systems continue to increase, management responsibilities of the storage systems can become distributed. That is, different users of a system will have different responsibilities with respect to
20 the objects stored within the system. In a storage system or storage array configured in this way, the storage management server facilitates the distribution of these responsibilities. In addition, a client can determine the configuration of the storage system, including the type, properties and status of storage connectivity devices (e.g.,
25 hubs, switches) and hosts, as well as, the version of software code or update used by the storage system. Moreover, distribution of the management of the storage systems is completely transparent to the clients.

30 With networks having numerous host computers connected to a large storage system (e.g., greater than 1 terabyte,) a tremendous number of input/output (I/O) requests are being made to and from the host computers at

any given time. For example, when used with such networks, the storage management server allows a client to determine that a large number of I/O requests are being made through one channel of the network. The client can then consider
5 the option of running an application using another storage system, thereby improving data throughput.

In another aspect of the invention, two or more storage management servers may be required to manage and distribute the information to clients. Such applications
10 include large storage applications having many storage devices and host computers. In such applications, each storage management server provides information relating to the operation status of the storage devices to the clients.

Embodiments of these aspects of the invention may
15 include one or more of the following features.

The storage management server includes a poller for gathering the information relating to the operation status of the storage devices. Among other responsibilities, the poller polls each of the storage devices at predetermined
20 intervals (as frequently as once per minute) to maintain the current status of the operation of each of the storage devices. Thus, status information for each storage device is guaranteed to be current or "fresh."

The storage management server includes an object
25 store, which serves as a central repository for storing the information relating to the operation status of the storage devices and an object server for distributing the information relating to the operation status of the storage devices to the clients.

30 The storage management server further includes a security component for limiting access by a client to the information stored at the storage management server. In this way, access privileges to one or more of the storage

devices can be limited to a particular user or groups of users. The storage management server further also includes a web server for facilitating communication with the clients, via the Intranet.

5 Each client further includes a graphical user interface for displaying the information relating to the operation status of the storage devices.

 The host computers and storage devices can be of different types. For example, one host computer could be a
10 UNIX-based system, while another is based on Windows NT. Similarly, for example, the storage systems can be different versions of EMC's Symmetrix storage systems.

 Where multiple management storage servers are used, the network architecture may further include a naming
15 service, connected to each of the storage management servers, to determine which of the central repositories of the storage management servers includes the information relating to the operation status of the storage devices of interest.

20 Another aspect of the invention relates to a method of managing a storage system including storage devices, the storage system communicating data to and from host computers. The method includes the following steps. A storage management server is provided between a number of
25 clients and the storage devices. The storage management server collects information relating to the configuration of the storage system and provides that information to at least one of the clients.

 Embodiments of this aspect of the invention may
30 include one or more of the following features.

 Providing information relating to the operation status of the storage devices includes using a poller to gather the information relating to the operation status of

the storage device. The information relating to the operation status of the storage devices is stored in a central repository on the storage management server. An object server distributes the information relating to the operation status of the storage devices to the clients.

A poller is used to poll each of the storage devices at predetermined intervals to maintain the current status of the operation of each of the storage devices.

In certain embodiments, information relating to the operation status of storage connectivity devices, which connect the hosts to the storage devices is also provided.

Other features and advantages of the invention will become apparent from the following description, including the claims and the drawings.

Brief Description of the Drawings

Fig. 1 is a block diagram of a network architecture of the invention.

Fig. 2 is a block diagram of the storage management server of the network architecture of Fig. 1.

Fig. 3 is an example of a display screen of the graphical user interface for accessing the storage management server of Fig. 1.

Fig. 4 is an alternative embodiment of a network architecture of the invention.

Detailed Description

Referring to Fig. 1, a network architecture 10 includes an enterprise disk storage 12 connected to host computers (hereafter, hosts) 14a-14e through a communications network 13. The term "enterprise" as used here means that the disk storage is configured to allow multiple connectivity by, for example, hosts provided by

different vendors. As shown here, disk storage system 12 includes individual storage arrays 12a-12e, which may be, for example, one of a number of different Symmetrix systems, products of EMC Corporation, Hopkinton, MA. Each of the
5 storage arrays 12a-12e typically has an excess of 1 gigabit of memory and is logically divided, in accordance with known techniques, into a plurality of logical volumes. Each storage array 12a-12e can thus have a plurality of logical
10 volumes, for example, four, eight, or more logical volumes in a single physical disk drive element. Details concerning the architecture and operation of these systems are found, for example, in Symmetrix Product Manuals for Models 5500, 52XX, 5100, 3500, 32XX, and 3100) all of which are incorporated herein by reference.

15 The communications network may be, for example, a Fibre channel network allowing any of host computers 14a-14c to communicate with any of storage arrays 12a-12e. Fibre channel network may be configured as a loop, a fabric having hubs or switches, or combinations of both. The hubs or
20 switches represent storage connectivity devices for forwarding data packets to appropriate ports based on the address of the packet.

Interconnecting the hosts 14a-14e and the storage arrays 12a-12e are disk drive controllers (not shown), for
25 example, that which is manufactured by EMC Corporation and known as the Symmetrix controller.

Network architecture 10 also includes a storage management server 16 which is connected to and communicates with each of the hosts 14a-14e via an appropriate agent 18a-
30 18e stored on the host. As is known in the art, an agent is a program that performs information gathering or processing in the background. In particular, each agent 18a-18e has the well-defined task of gathering and providing information

to management server 16 associated with the operation status of the storage array associated with the host running the agent. As will be described in greater detail below, management server 16 stores and makes accessible to a number of clients 20a-20e the operation status information of the storage arrays. Clients 20a-20e communicate with management server 16 through an on-line object request broker 22 (e.g., Voyager ORB, a product of ObjectSpace, Inc., Dallas, TX). Object request broker 22 serves as a remote access adapter interface for ensuring that a client is able to access management server 16. Other object request brokers 22 including RMI and CORBA are equally suitable for use for linking a client to management server 16. Thus, clients 20a-20e can monitor the status of any or all storage arrays 14a-14e from virtually anywhere in the world, or at least where Intranet access is available. In particular, each of clients can independently access management server and determine the status of any or a number of storage systems 12a-12e. For example, a client 20a can determine whether throughput problems exist between host 14b and storage arrays 12b and 12c. Further, client 20a can determine the configuration of storage array 12e including the version of software code or update used by that storage system.

Referring to Fig. 2, storage management server 16 includes a poller 30 which has the primary function of retrieving the pertinent information about each of storage arrays 12a-12e. Upon initiation of management server 16, poller 30 will poll for and "discover" all connected hosts, storage arrays, and storage connectivity devices (e.g., Fibre Channel, hubs, switches). Poller 30 will then begin monitoring basic status of the arrays, including how the storage arrays are defined and grouped, as well the status

of communication with agents 18a-18e. The frequency of polling can be selected by the client or, if not defined, be set to a default value (e.g., every minute.) The discovery process should also include retrieval of the storage array
5 configuration information.

Once retrieved, the configuration information is stored (in object format) in Object Store 32 and poller 30 polls the appropriate agent for a configuration change date and time. Upon recognizing a configuration change, poller
10 30 will retrieve a current or "fresh" copy of the configuration. Poller 30 will also notify the Object Server 34 (here, a Java Object Server), which is responsible for notifying registered clients of the new configuration. Of course, whenever poller 30 requests a status update of
15 configuration, it will poll for additional agents and storage arrays which may have been recently added.

Object Store 32 maintains persistence on its stored objects. Persistence is maintained for current status and configuration information for each storage array connected
20 to the network. In addition, statistical information may be stored for all or selected ones of the storage arrays.

Object Server 34 uses a security component 36 to first validate user ID codes and passwords from a client, and will then retrieve from Object Store 32, the current
25 status of all storage arrays that the user is authorized to view. Object Server 34 updates the client with any new data until the client "unregisters."

Storage management server 16 can optionally include an HTTP server (web server) 38, which is used to establish
30 initial connection with clients 20a-20e. In essence, HTTP server 38 allows clients to establish connection using the address of the HTTP server rather than the proxy (e.g., Voyager) address.

Referring to Fig. 3, an example of an operator display screen 40 of the graphical user interface is shown. Display screen 40 includes a Java applet window 42 contained within a web-browser window 44 (here, NetScape). Within
5 applet window 42 are several sub-windows, including a site sub-window 46 which lists the storage arrays accessible by the user, while a session sub-window 48 provides the status information of the current session. Sub-window 50 provides a graphical representation showing the grouping and
10 hierarchy of the storage arrays, so that the user is able to quickly understand the connectivity relationship of the storage arrays of interest.

Referring to Fig. 4, in another embodiment, two or more storage management servers 16a, 16b can be used to
15 distribute the monitoring load of the storage arrays. Each management server includes the structure and provides the functionality shown in Fig. 2 as described above. In these applications, however, because multiple management servers are used, a name server 60 is provided to help manage the
20 distribution of the management resources. In particular, name server 60 determines which of the Object Stores of storage management servers 16a, 16b includes the information relating to the operation status of the storage arrays of interest. Note that from the viewpoint of clients 20a-20e,
25 the use of multiple management servers is invisible.

Other embodiments are within the scope of the claims. For example, in the above description, the storage management server is used in conjunction with a Symmetrix enterprise disk storage system. It is important to
30 appreciate that the storage management server is equally applicable in use with other storage area networks and other storage devices, including those that are not Symmetrix devices.